

RSC Basis Platform

Microagent Platform for Cluster Management

Pavel Lavrenko, Mikhail Malkov

11 September 2018

Agenda

- BasIS Automation Platform
- BasIS HPC Stack
- Live Demo

About RSC



RSC is a supercomputer company

Top500 Vendors System Share List (2014)
installations in the Russian Top50

9th place
24% (12)

Percentage of liquid-cooled components
Leading compute density
Worldwide record in energy density
The most energy efficient system in Russia

100%
153 nodes per rack
400 kW per rack
PUE 1.027

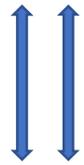
Key Products



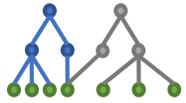
- Liquid-cooled high-performance compute platform
- Compute servers
- High-performance storage system
- **End-to-end datacenter management platform**



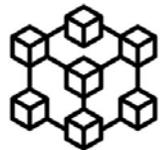
RSC Basis Platform



Vertical integration of Hardware, Software and Infrastructure



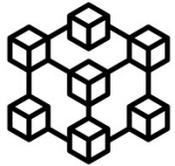
Knowledge about all datacenter objects and their relationships



Microagent Mesh for Cluster Automation



Microagent Framework for Datacenter Automation



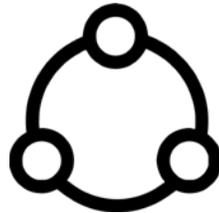
Distributed
management apps



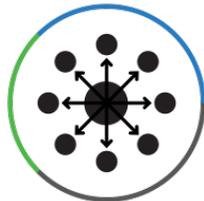
Agents



App Repository



Agent Lifecycle



Messaging system



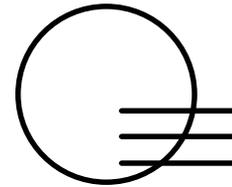
SDK



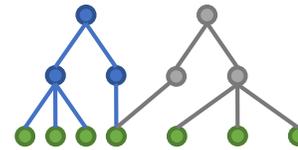
CMDB Cluster Management Database



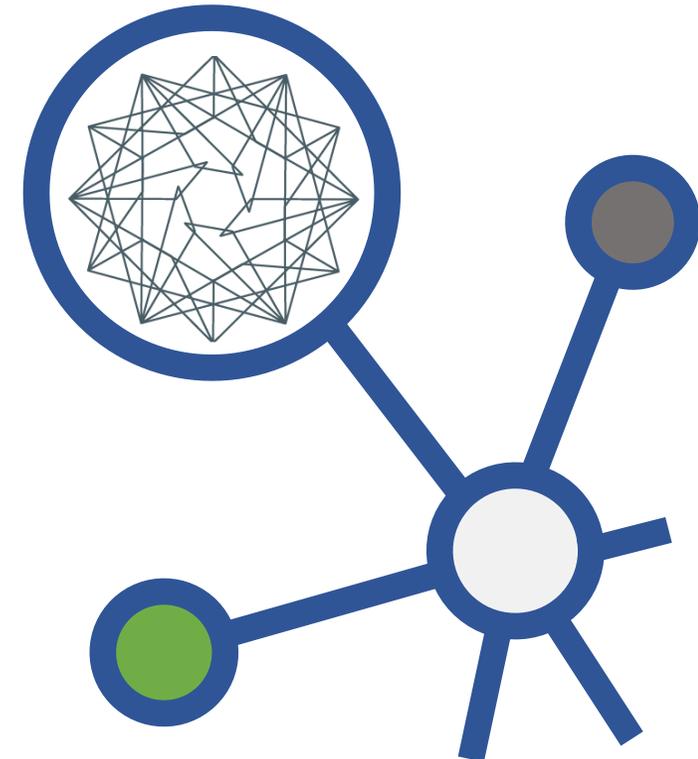
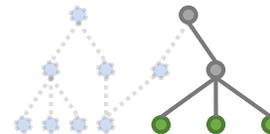
All objects and its attributes



All topologies



Query Language (QDSL)



BasIS HPC Stack

Based on BasIS Automation Framework



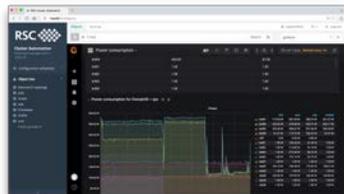
Advanced node management: discovery, provisioning and control



Powerful Automatic Configuration Engine



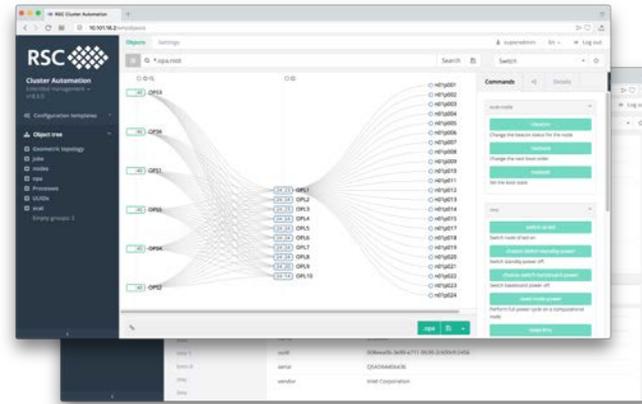
Scheduler management and advanced accounting



Monitoring



Advanced tools



RSC Basis HPC Stack Demo

Mikhail Malkov

Check our Poster for more details!

In Posters Section

Storage-on-demand

RSC Tornado hyper-converged solution for data processing

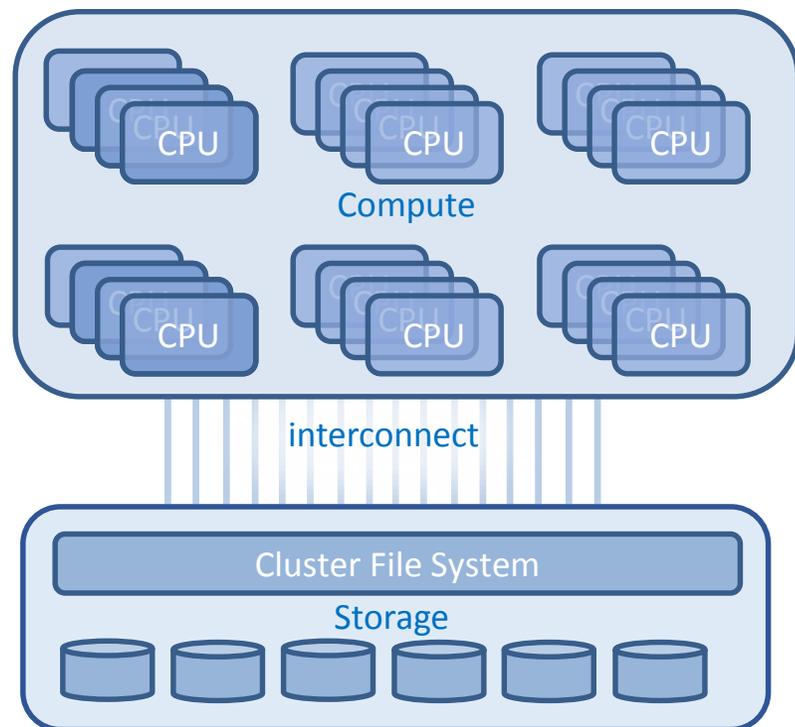
Pavel Lavrenko

11 September 2018

Generations of Storage Architecture

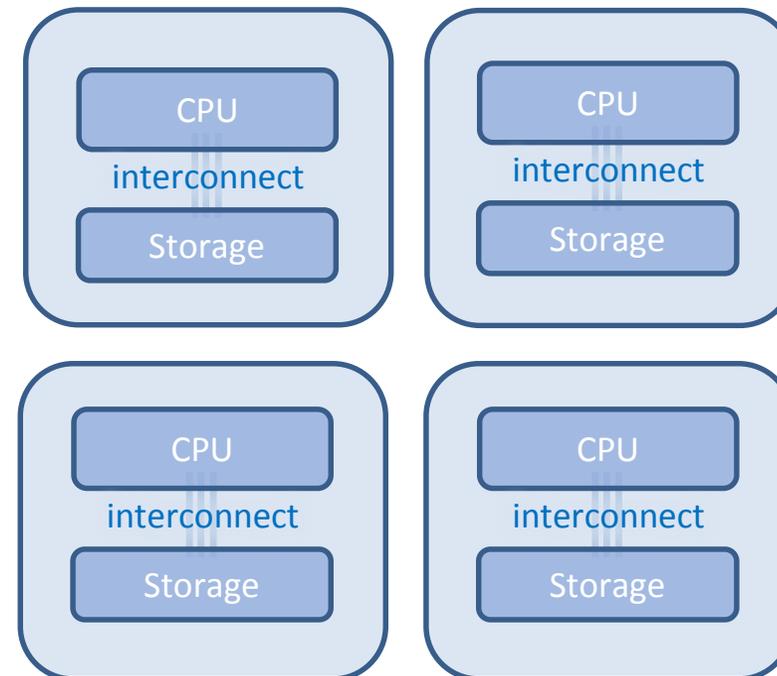
Standalone NAS

Based on dedicated storage hardware



Converged storage

Uniting storage with compute in dedicated appliances



Current Storage Architectures for Clusters

- Standalone or Clustered
- Hardware-based or Software-based redundancy
- Block, Object or File

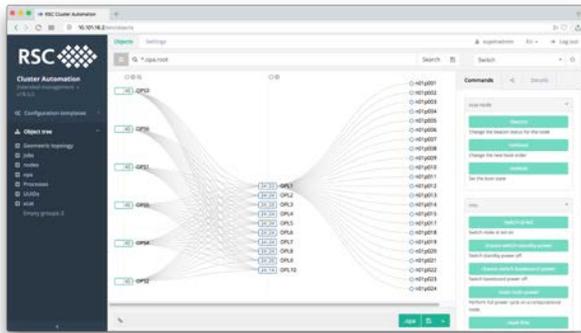
Why are modern solutions not universal?

They are designed to solve particular demand or to be optimized by performance, capacity or the cost of hardware, software or support

All modern solutions are united by same limitations:

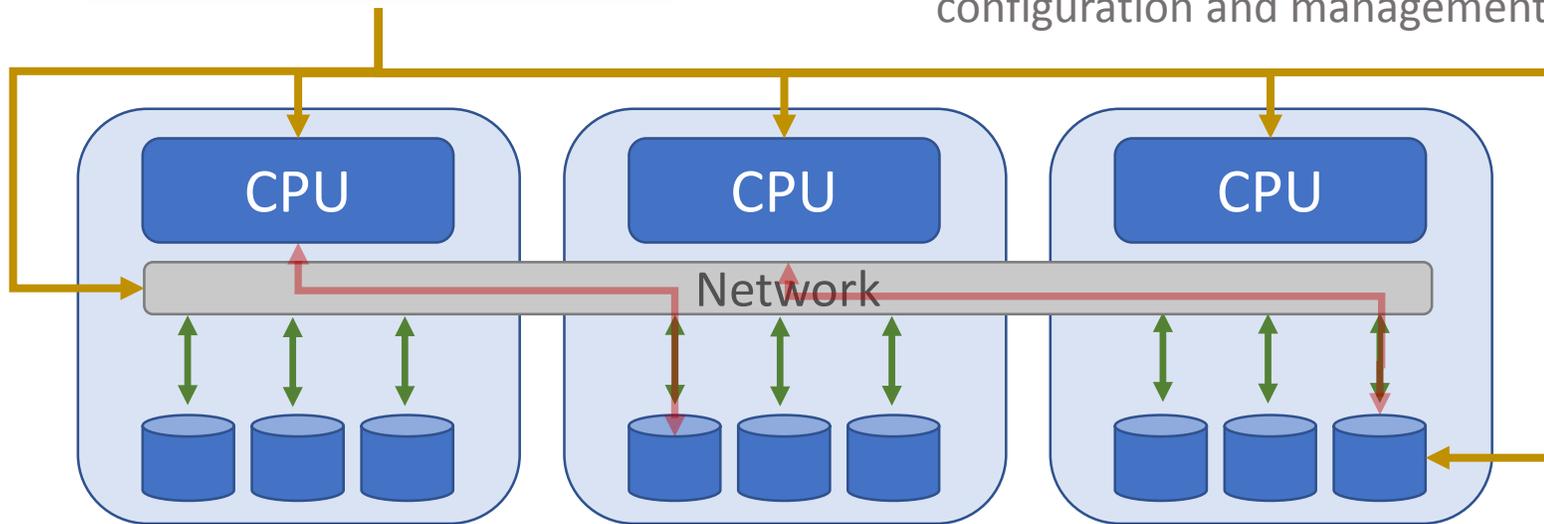
- Storage systems for clusters are always separate part of equipment
- A particular storage solution has an architecture, fixed at installation – topology, configuration, access protocols and etc.

Hyperconverged storage



Software Defined Storage
Integrated with full software stack for
cluster management

configuration and management



Hyperconverged storage

Is a Storage System on Demand

It can create dedicated filesystem on demand, even for a single job run

Works with any type of Distributed File System

Lustre, BeeFS, EOS, GPFS and etc.

Based on Unified Hardware Components

Each node can act as compute, storage or both. No need for dedicated storage hardware

Has Linear Scalability

Scaling is achieved by adding more nodes to a cluster

Is a Software Defined Storage

It is based on software defined configuration fully optimized for the user needs

Components of RSC's hyperconverged liquid cooling platform



Unified rack

Completely liquid-cooled
Redundant power
Up to 153 modules
Fast 100 GBPs network



Compute & storage nodes

2 CPUs / 384 GB RAM
Up to 12 NVMe drives
Hot-swappable



Standard compute node

Ultra high density



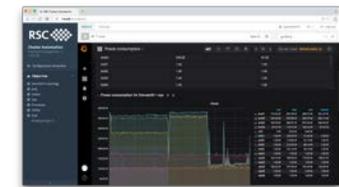
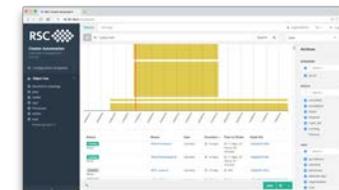
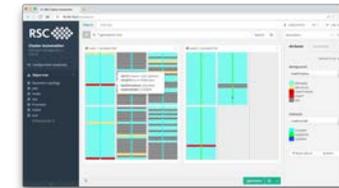
Fat compute node

several operational modes



Microagent Datacenter Automation Platform

Runs on-demand storage
on the fly over the network



New Compute & Storage Node



2 x Intel Xeon Skylake-SP processors with 28 cores each

768 GiB DDR4 2666

2 x Omni-Path 100 Gb/s or or EDR InfiniBand

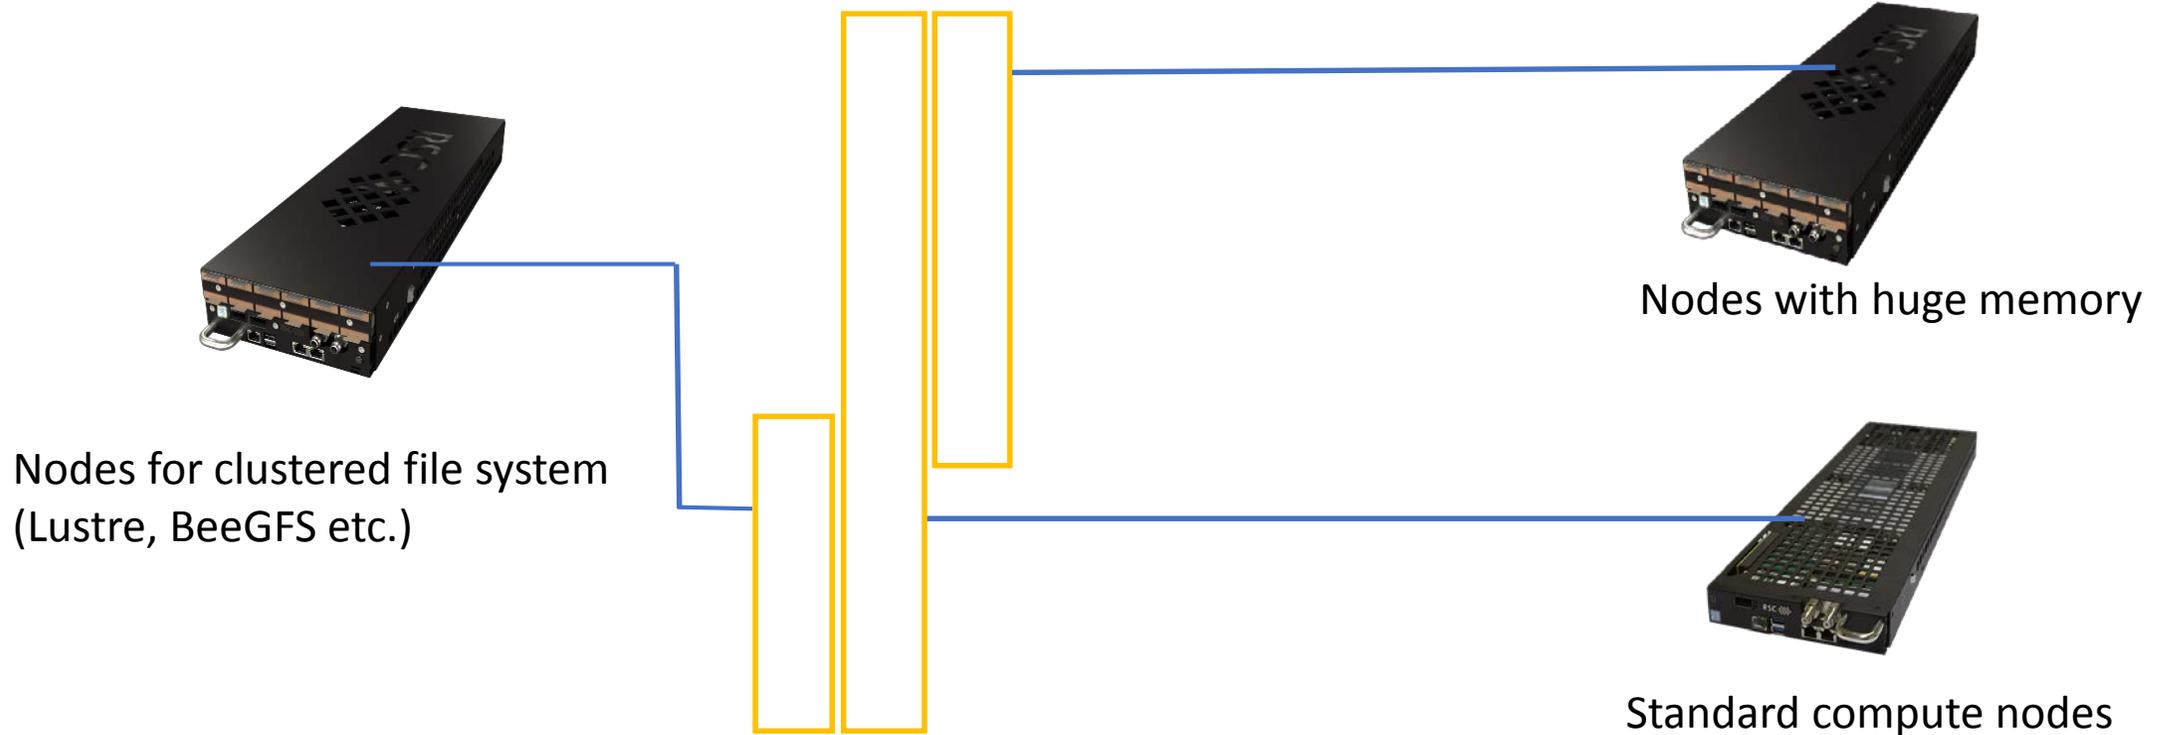
12 hot-swap NVMe SSDs

100% 'hot water' liquid cooled solution

Intel® SSD DC P4511 NVMe

Intel® Optane™ SSD DC 4801X (M.2)

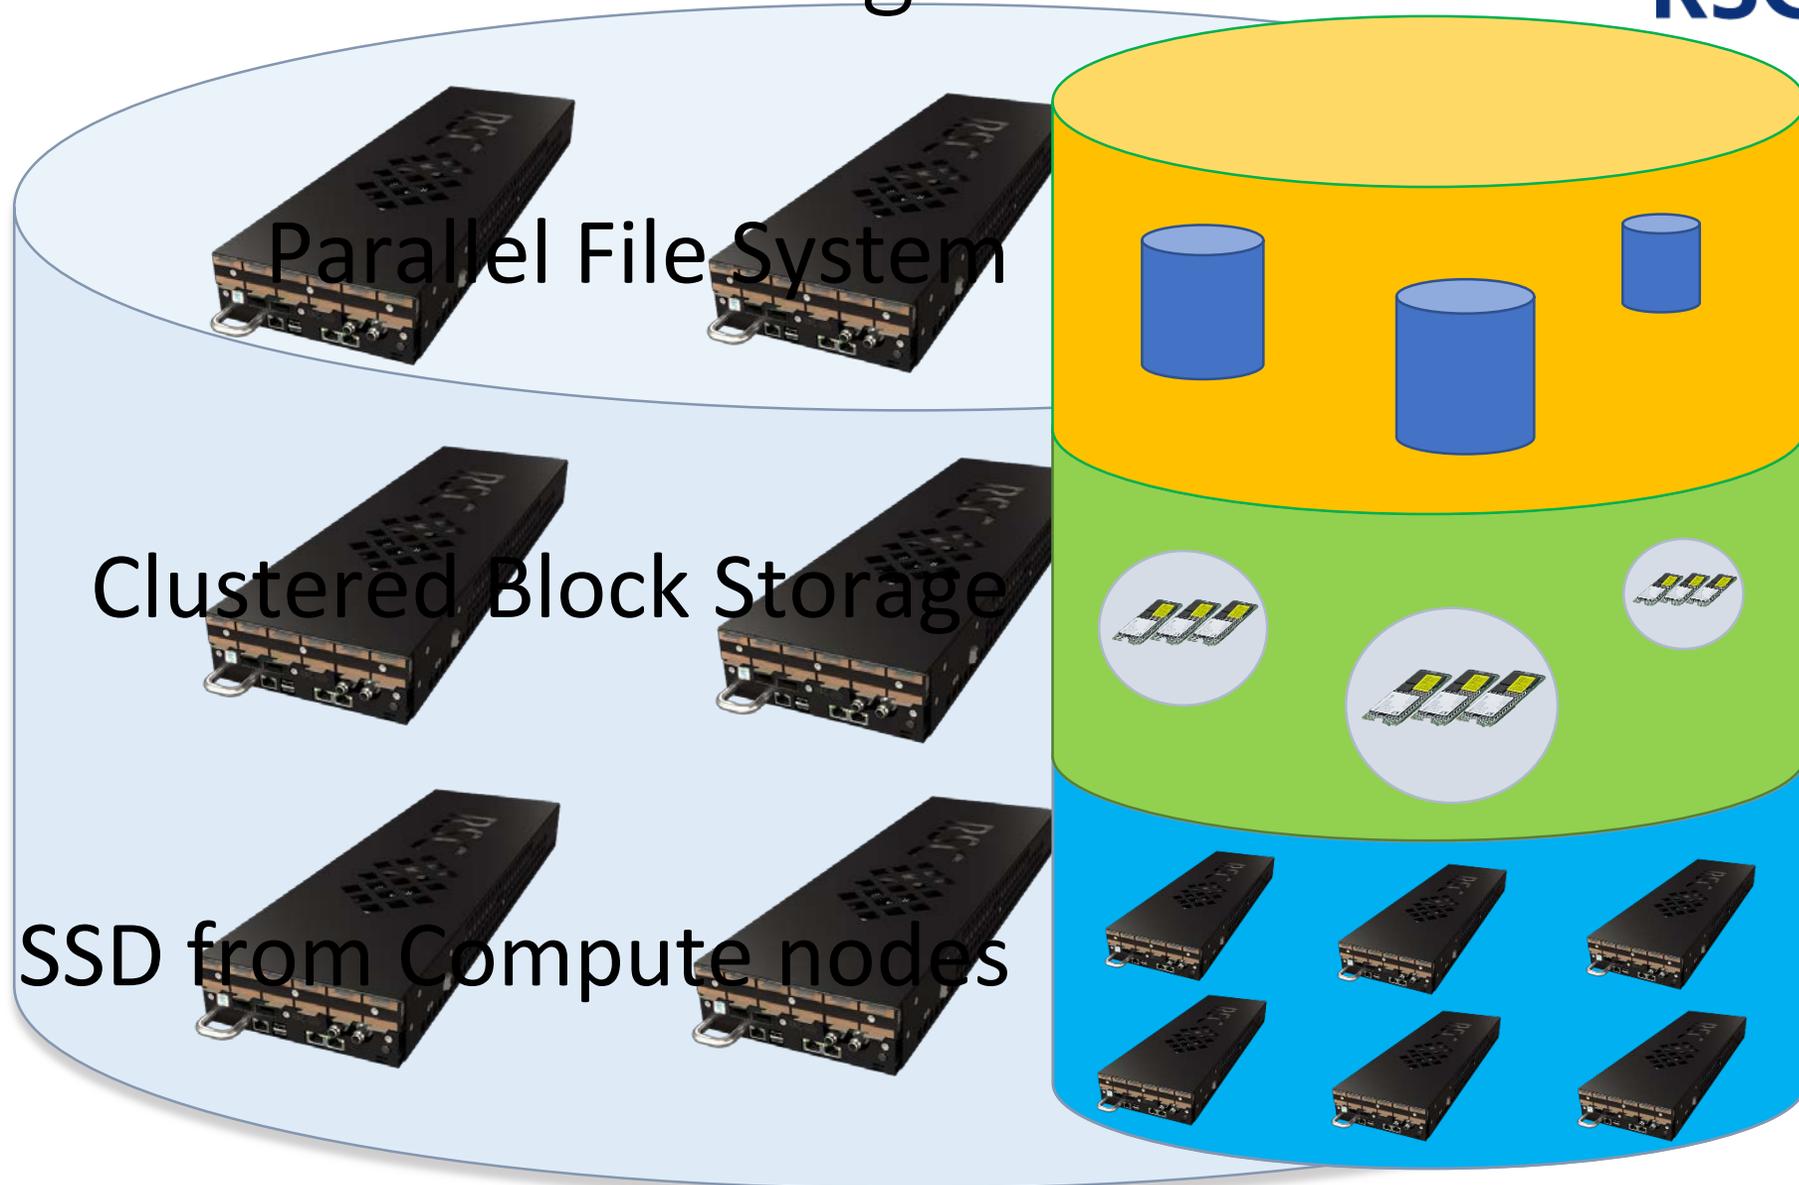
Variety of Hyperconverged node types



Compute and Store

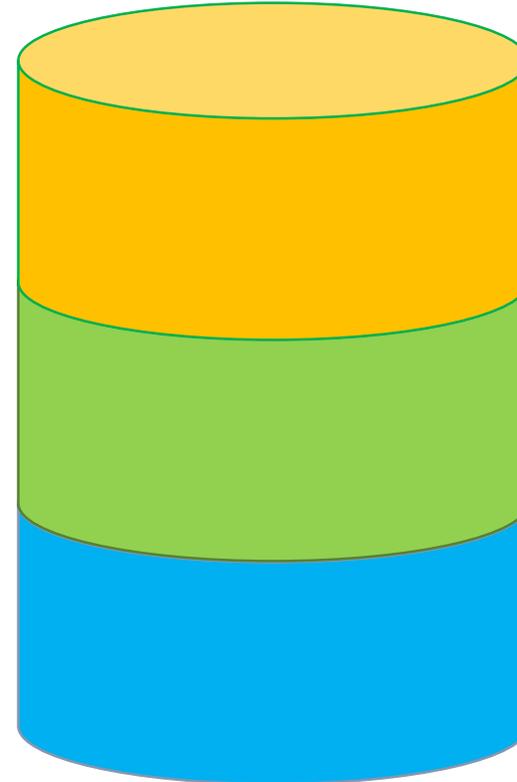


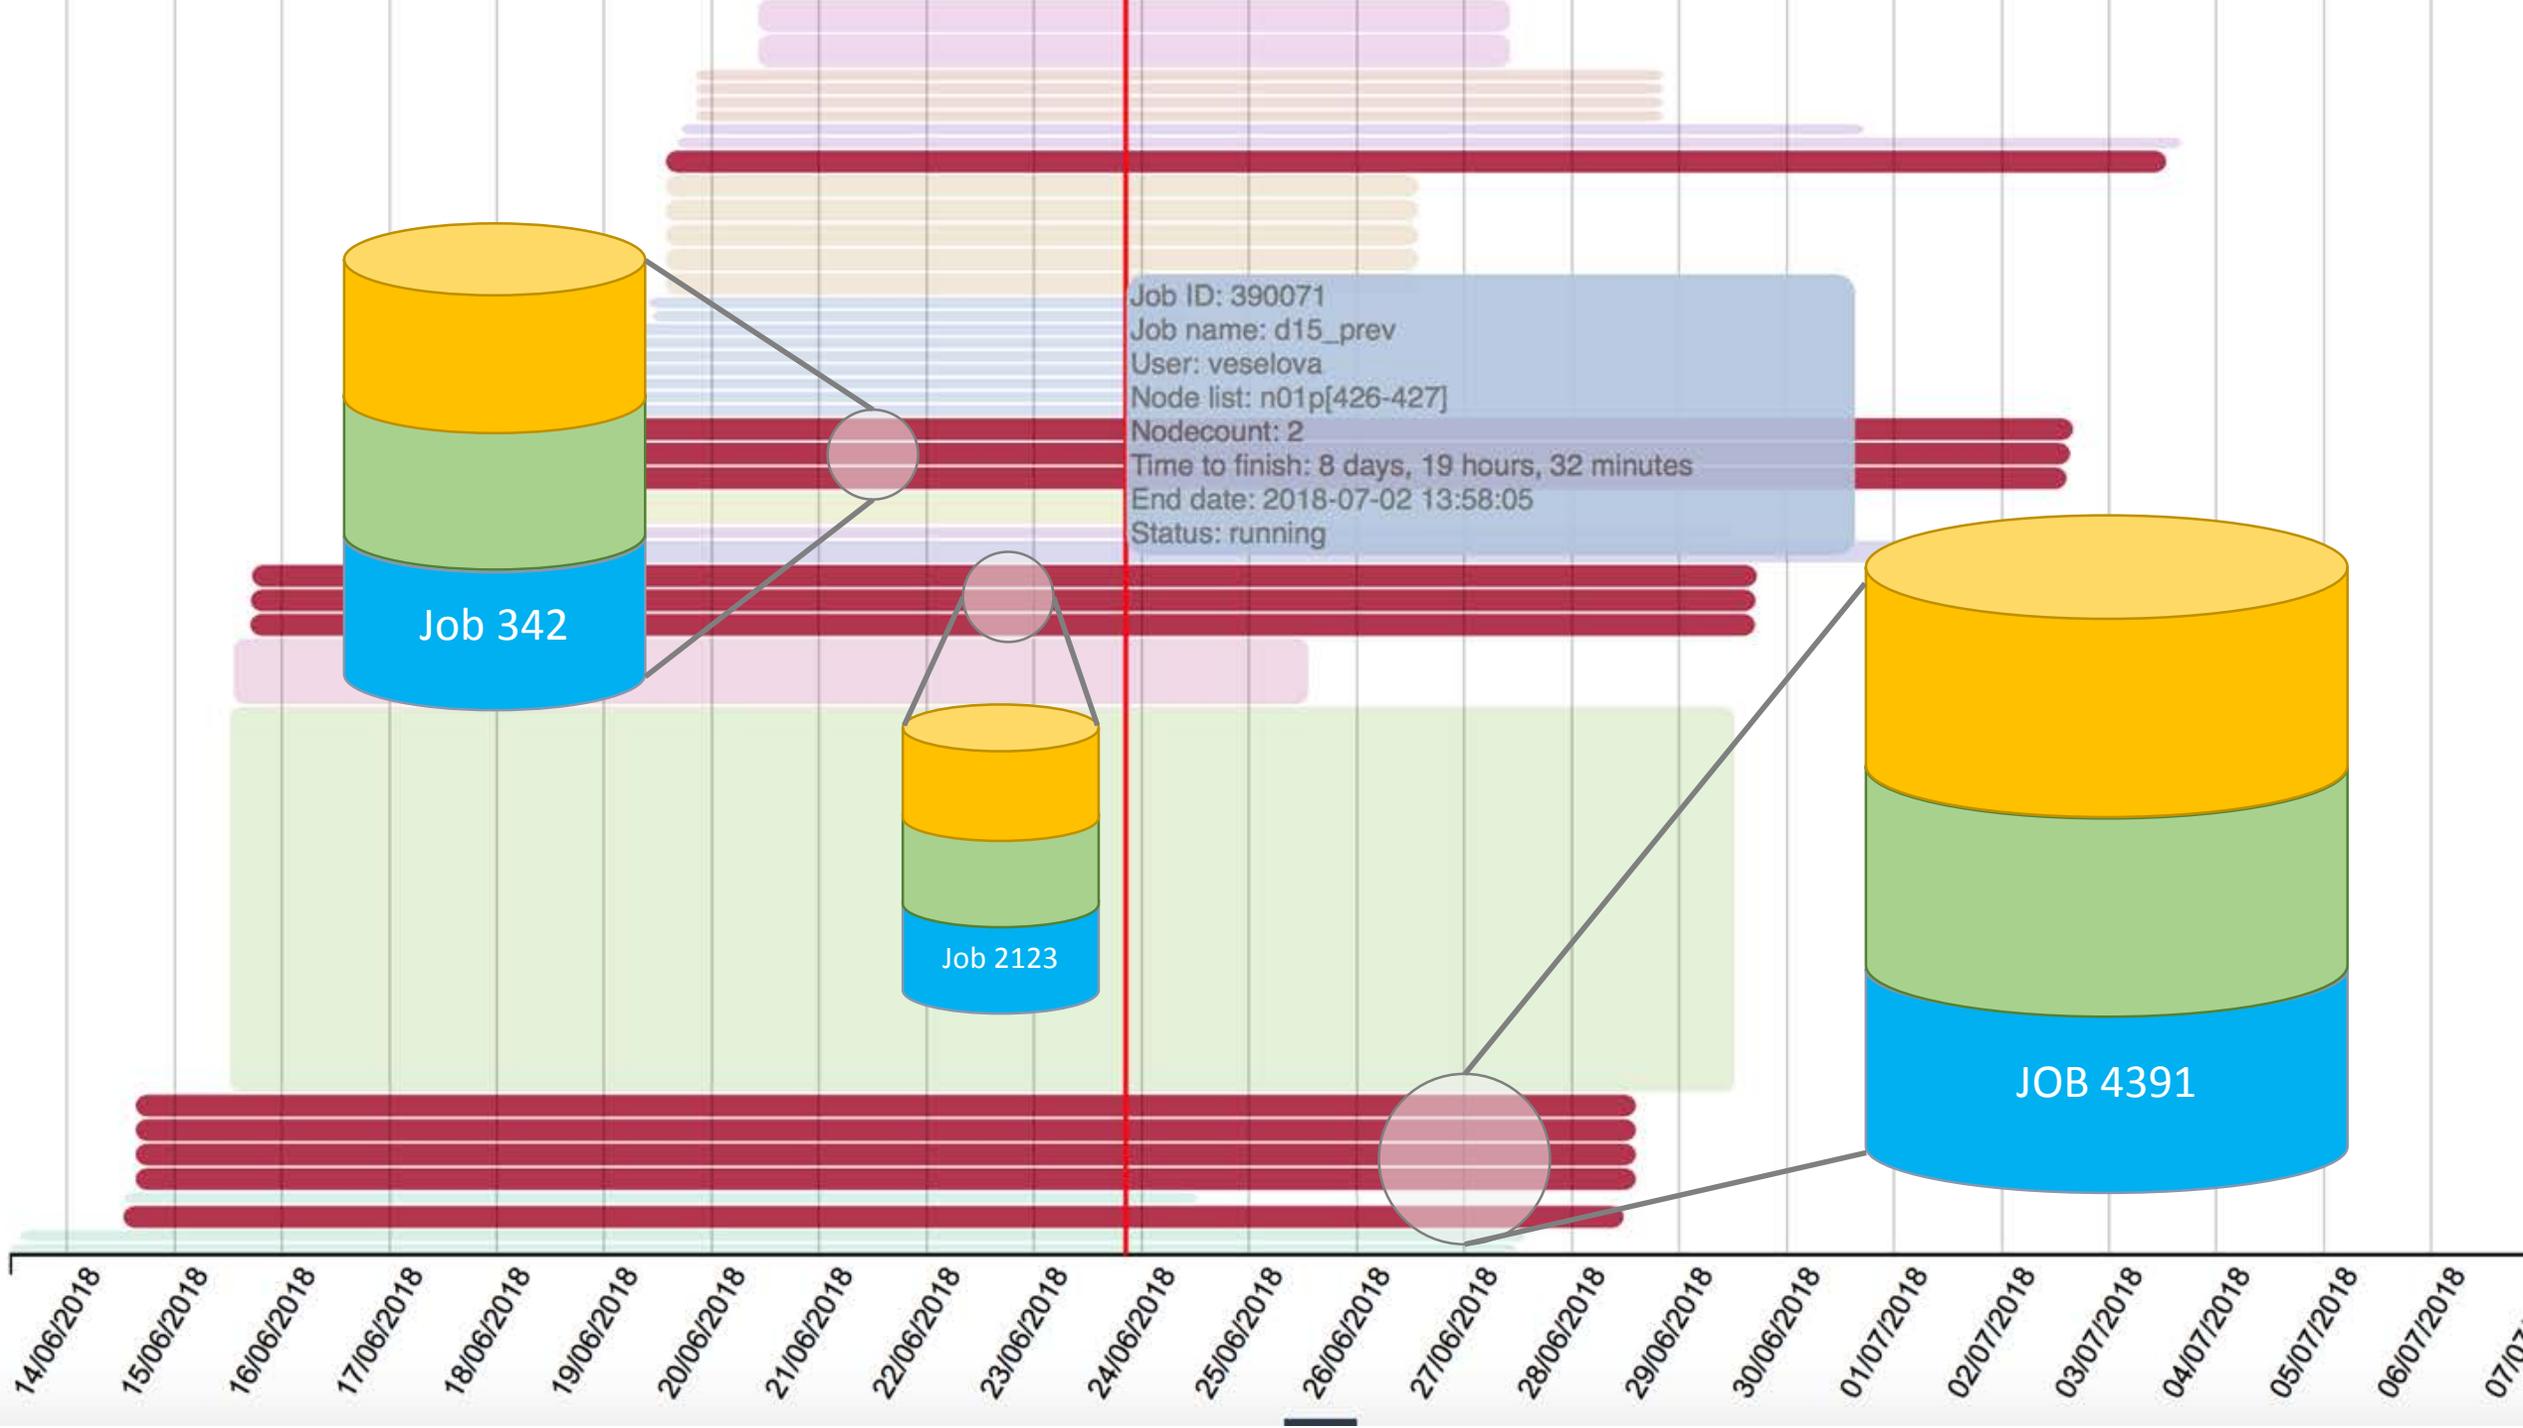
Software Defined Storage



Features

- Object or filesystem operational mode
- Different file systems to provision
- Storage or memory operational mode
- Adjustable level of redundancy
- Configurable lifetime
- Adjustable storage capacity



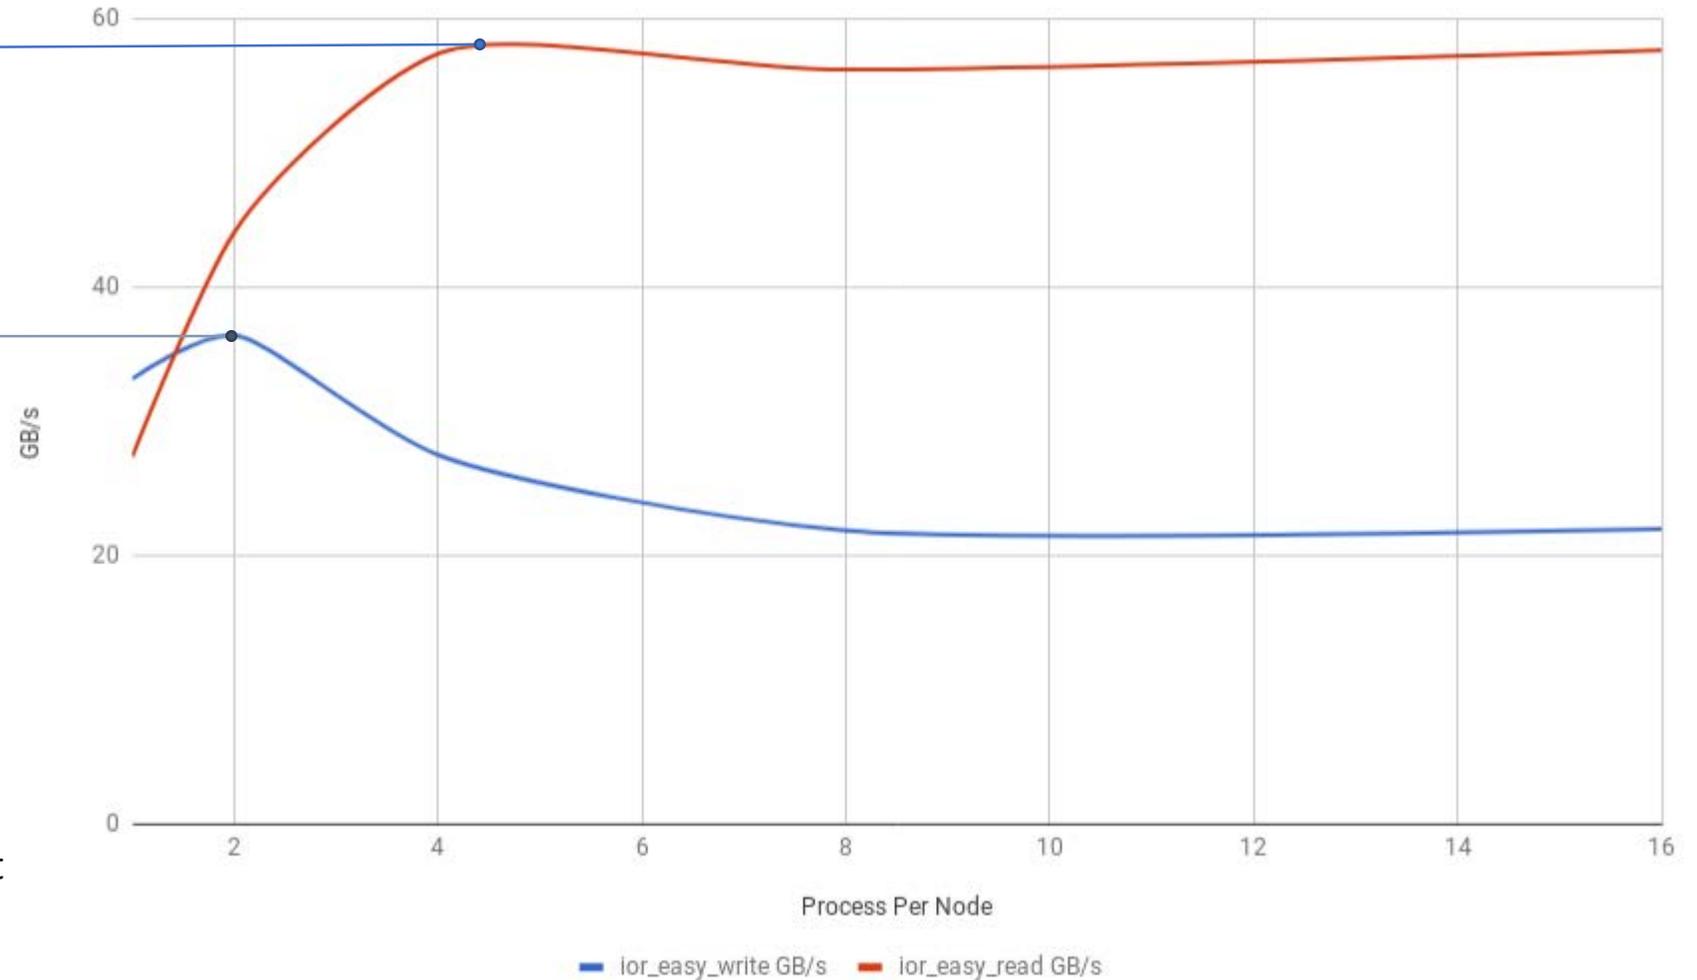


Lustre Storage-on-Demand io500 benchmark results



56 GB/s easy_read ior test

36 GB/s easy_write ior test



Configuration:

1 x MDS

12 x OSS (6 x 1TB NVMe SSD)

24 Lustre Clients

100Gbps Intel Omni-Path interconnect

9th in io500 benchmark



IO-500

This is the official *ranked* list from [ISC-HPC 2018](#). The list shows the best result for every given combination of system/institution/filesystem (i.e. multiple submissions from the same system are not shown; only the most recent is shown). The full list is available [here](#).

| # | information | | | | | | io500 | | |
|----|----------------|--------------------------------------|----------------|----------------|--------------|------|--------|--------|--------|
| | system | institution | filesystem | storage vendor | client nodes | data | score | bw | md |
| | | | | | | | | GiB/s | kIOP/s |
| 1 | Oakforest-PACS | JCAHPC | IME | DDN | 2048 | zip | 137.78 | 560.10 | 33.89 |
| 2 | ShaheenII | KAUST | DataWarp | Cray | 1024 | zip | 77.37 | 496.81 | 12.05 |
| 3 | ShaheenII | KAUST | Lustre | Cray | 1000 | | 41.00* | 54.17 | 31.03* |
| 4 | JURON | JSC | BeeGFS | ThinkparQ | 8 | | 35.77* | 14.24 | 89.81* |
| 5 | Mistral | DKRZ | Lustre2 | Seagate | 100 | | 32.15 | 22.77 | 45.39 |
| 6 | Sonasad | IBM | Spectrum Scale | IBM | 10 | zip | 24.24 | 4.57 | 128.61 |
| 7 | Seislab | Fraunhofer | BeeGFS | ThinkparQ | 24 | | 16.96 | 5.13 | 56.14 |
| 8 | Mistral | DKRZ | Lustre1 | Seagate | 100 | zip | 15.47 | 12.68 | 18.88 |
| 9 | Govorun | Joint Institute for Nuclear Research | Lustre | RSC | 24 | zip | 12.08 | 3.34 | 43.65 |
| 10 | EMSL Cascade | PNNL | Lustre | | 126 | | 11.12 | 4.88 | 25.33 |
| 11 | Serrano | SNL | Spectrum Scale | IBM | 16 | | 4.25* | 0.65 | 27.98* |
| 12 | Jasmin/Lotus | STFC | NFS | Purestorage | 64 | zip | 2.33 | 0.26 | 20.93 |



RSC Storage-on-Demand provides outstanding performance at $1/20^{\text{th}}$ of traditional parallel file system cost

Check our Poster for more details!